

Formal Modelling, Testing and Verification of HSA Memory Models using Event-B

Ashish Darbari
and Iain Singleton
Imagination Technologies
Kings Langley
Hertfordshire WD4 8LZ, UK

Michael Butler
and John Colley
University of Southampton
Southampton
Hampshire SO17 1BJ, UK

Abstract—The Heterogeneous System Architecture (HSA) Foundation has produced the HSA Platform System Architecture Specification that goes a long way towards addressing the need for a clear and consistent method for specifying weakly consistent memory. A weakly-consistent memory model is one of the fundamental cornerstones for achieving high performance concurrency with low power on mobile platforms. HSA is specified using a natural language which makes it open to multiple ambiguous interpretations and could thereby render bugs in implementations of it in hardware and software. In this paper, we present a formal model of HSA which can be used in the development and verification of both concurrent software applications as well as in the development and verification of the HSA-compliant platform itself. We use the Event-B language to build a provably correct hierarchy of models from the most abstract to a detailed refinement of HSA that is close to implementation level. Our memory models are general in the sense that they represent arbitrary numbers of masters, programs and instruction interleavings and we reason about such general models using theorem proving. By using the Rodin tool for Event-B we are able to seamlessly model and verify the entire hierarchy of models using proofs to establish that each refinement is correct. We also define an automated validation method that allows us to check baseline compliance of the model against a suite of published HSA litmus test cases. Once we have completed the model validation we develop a coverage driven method to extract a richer set of test cases from the formal Event-B model and a user specified coverage model. These tests are then used for extensive regression testing of hardware and software systems. We believe our methodology of refinement based formal modelling, baseline compliance testing of the formal model and coverage driven test extraction using a single language of Event-B language and the Rodin tool is a completely new way of addressing a profoundly important challenge facing the design and verification of low-power of multi-core systems.

I. INTRODUCTION

Weakly-consistent memory [1] is one of the fundamental cornerstones for achieving high performance concurrency with low power on mobile platforms. Although the buffering mechanisms that underly weakly consistent memory models are relatively straightforward to understand, developing a high-level, natural language specification that represents the behaviours of all implementations is very difficult. This presents a challenge to both software engineers, who wish to develop efficient, race-free concurrent programs and for platform engineers who wish to develop systems that conform to a particular memory model.

The Heterogeneous System Architecture (HSA) Foundation is a not-for-profit industry standards body which has produced the HSA Platform System Architecture Specification [2] that goes a long way towards addressing the need for a clear and consistent method for specifying weakly consistent memory, but since it is still in natural language there is still a need for a more formal representation. In this paper, we describe how the HSA specification and the well-founded terminology and concepts it describes, is used to develop a formal model in Event-B [3] which can be used in the development and verification of both concurrent software applications as well as in the development and verification of the HSA-compliant platform itself, the specific configuration of masters, slaves and interconnect.

We use the Event-B [3] language to build a provably correct hierarchy of models from the most abstract to a detailed refinement of HSA that is close to implementation level. Our memory models are general in the sense that they represent arbitrary numbers of masters, programs and instruction interleavings and we reason about such general models using theorem proving. By using the Rodin tool for Event-B we are able to seamlessly model and verify the entire hierarchy of models using proofs to establish that each refinement is correct. We also define an automated validation method that allows us check baseline compliance of the model against a suite of published HSA litmus test cases. The validation method involves refining the general models to represent litmus tests and uses model checking to verify that all possible interleaving of a litmus test achieve the expected outcome. Once we have completed the model validation we develop a coverage driven method to extract a richer set of test cases from the formal Event-B model and a user specified coverage model. These tests are then used for extensive regression testing of hardware and software systems. We believe our methodology of refinement based formal modelling, baseline compliance testing of the formal model and coverage driven test extraction using a single language of Event-B language and the Rodin tool is a completely new way of addressing a profoundly important challenge facing the design and verification of low-power of multi-core systems.

II. EVENT-B BASED FORMAL MODELLING OF HSA

Event-B is a proof-based modelling language and method that enables the systematic development of specifications using a formal notion of refinement. The Rodin platform [3] is the Eclipse-based IDE that provides automated support for Event-B modelling, refinement and mathematical proof, model-checking and model-based test generation.

Our approach starts at the abstract level, focusing on issue and observation of memory instructions. In subsequent refinements, we introduce *set-theoretic relations* on the memory accesses which constrain the order in which they can be issued or observed. At each refinement level we are able to verify properties formally at their appropriate level of abstraction and detailed instruction sets are introduced in correctness-preserving refinement steps. This refinement-based approach helps to manage complexity of the modelling and verification.

Using this modelling approach, we build a formal memory model that represents an arbitrary set of masters running an arbitrary set of program threads. We then use further Event-B refinements to constrain the model to represent a fixed set of masters, each executing a specific program thread. In this way, we can represent a published *HSA litmus test*, together with a property or set of properties which specify the desired *outcome* of the test. We then use a combination of *automatic theorem proving, model checking and simulation* to verify and validate the formalised memory models against the formalised litmus tests. Once we have completed the model validation we develop a coverage driven method to extract a richer set of test cases from the formal Event-B model and a user specified coverage model. These tests can then be used for extensive regression testing of hardware and software systems.

In Figure 1 below we show the Event-B model refinement hierarchy. The abstract generic models are denoted by GM_n , where n ranges from 1 to 3 whilst the HSA-specific models are denoted by $HSAM_n[_m]$ where n ranges from 4 upwards and $_m$ represents a litmus test at the concrete level. The abstract models, $GM1$ to $GM3$, are generic and can form the foundation for modelling, for instance, Sequential Consistency (SC) [4], Total Store Order (TSO) [5] or ARM [6] memory models. In $GM1$ we define the fundamental memory accesses with the events *Issue* and *Observe*, in $GM2$ we differentiate between *LOAD* and *STORE* memory accesses and in $GM3$ we model the *observation* of the *LOAD* and *STORE* values. In the next refinement, $HSAM4$, we model the ordering rules for HSA atomics and FENCES and in refinement $HSAM5$ we model the registers generically. We then introduce a series of refinements, $HSAM6[1 .. m]$ to model a series of m concurrent programs, the litmus tests, each running on a fixed set of masters and use the ProB model checker [7] to validate the model. Each of these models is then further refined, $HSAM7[1 .. m]$ to introduce the functional coverage metric for coverage measurement and test generation.

Once we have completed the model validation and have a set of litmus tests which cover the model, we relax the constraints on the litmus test models $HSAM7[1 .. m]$ so that instead of




Fig. 1. Model Refinement Hierarchy

each representing a single, fixed concurrent program, each model represents a class of concurrent, synchronised programs, based on the litmus test, from which we can generate tests automatically that can be used in the development and verification of the HSA-compliant platform itself.

In summary, the contributions of this paper are:

- a hierarchy of formalised and verified memory models with varying ordering constraints
- a method for validating models against known litmus tests through model-checking
- a method for automatic generation of tests from a constrained model for regression testing of an HSA-compliant platform.

III. UNDERSTANDING HSA ORDERING

We develop a generalised approach to weakly consistent memory modelling, based on HSA terminology, which provides the building blocks with which we can model weak memory semantics with a single, consistent modelling framework. We also provide at the concrete level a generalised notation for representing litmus tests and the invariants associated with those tests.

We begin with a definition of *load* and *store* from the HSA reference manual [2]. We focus our modelling around *LOAD* and *STORE* operations that may be issued and observed by different units of execution.

- For primitives of type *store*, visibility to unit of execution A of a store operation X is when the data of store X is available to *loads* from unit of execution A.
- For primitives of type *load*, visibility is when a load gets the data that the unit of execution will put in the *register*.

We shall use the HSA definition as the basis for formalising the specification of *loads* and *stores*.

| P1 | P2 | P3 |
|---------------|----------------|----------------|
| p1i1: S 1, a1 | p2i1: L R1, a2 | p3i1: L R1, a1 |
| p1i2: S 1, a2 | p2i2: L R2, a1 | p3i2: L R2, a2 |

Fig. 2. Litmus test

A. Observation Ordering

Without considering synchronisation, based on the HSA definition above, we define observation ordering, using the litmus test of Figure 2 to illustrate the definitions. We assume that the values at addresses *a1* and *a2* are initially set to 0

- A STORE is observed by any master when a subsequent LOAD issued by that master would result in the LOAD returning the value associated with the STORE
 - The STOREs *p1i1* and *p1i2* are observed by masters *P1*, *P2*, *P3* in some order and these STORE observations with respect to the LOAD observations determine the value which the LOADs observe.
- A LOAD is observed by the master issuing the LOAD when the LOAD returns the value associated with the last observed STORE to the LOAD's address
 - The LOADs *p2i1* and *p2i2* are observed by *P2* in some order; similar for the *P3* LOADS.

Note that this definition of observation ordering does not refer to values in registers. We use, instead, the more abstract concept of a LOAD *returning* a value. We also use the term *observe* [8] rather than the HSA *visibility*.

For *Synchronisation*, the ordering of STORE and LOAD observations is constrained

- 1) An ordering is established between LOAD and STORE observations when those LOADs and STOREs are issued by a single master (e.g., program order). In Figure 3, the arrow indicates the ordering of the LOADs on *P2* and *P3*.

| P1 | P2 | P3 |
|---------------|----------------|----------------|
| p1i1: S 1, a1 | p2i1: L R1, a2 | p3i1: L R1, a1 |
| p1i2: S 1, a2 | p2i2: L R2, a1 | p3i2: L R2, a2 |

Fig. 3. Ordering between LOADs on the the same master

- 2) An ordering is established between LOAD observations on different masters to the same address as shown in Figure 4.

| P1 | P2 | P3 |
|---------------|----------------|----------------|
| p1i1: S 1, a1 | p2i1: L R1, a2 | p3i1: L R1, a1 |
| p1i2: S 1, a2 | p2i2: L R2, a1 | p3i2: L R2, a2 |

Fig. 4. Ordering between LOADs on different masters

- 3) An ordering is established between a STORE observation by a given master and an observation of a LOAD by a different master for the same address as shown by the

labelled arrow in Figure 5. *P2* observes the STORE *p1i2* before it observes the LOAD *p2i1*.

| P1 | P2 | P3 |
|---------------|----------------|----------------|
| p1i1: S 1, a1 | p2i1: L R1, a2 | p3i1: L R1, a1 |
| p1i2: S 1, a2 | p2i2: L R2, a1 | p3i2: L R2, a2 |

Fig. 5. Ordering between STOREs and LOADs on different masters

With these three orderings in place, the following outcome is not possible:-

(*p2*:R1 = 1 and *p2*:R2 = 0 and *p3*:R1 = 1 and *p3*:R2 = 0)

For this outcome to occur would mean that

- *p2i1* and *p2i2* are observed in program order
- *P2* must observe the LOAD *p2i2* before *P3* observes the LOAD *p3i1*
- *P3* must observe the LOAD *p3i2* before *P2* observes the LOAD *p2i1*
- This ordering is invalid (cyclic), therefore the assumed outcome is not possible, as shown in Figure 6.

| P1 | P2 | P3 |
|---------------|----------------|----------------|
| p1i1: S 1, a1 | p2i1: L R1, a2 | p3i1: L R1, a1 |
| p1i2: S 1, a2 | p2i2: L R2, a1 | p3i2: L R2, a2 |

Fig. 6. Invalid ordering

We now describe these three synchronisation orderings, using the definitions of *Program Order* (*po*), *Coherent Order* (*co*) and *Happens Before Order* (*hb*), as defined in the HSA specification. For *po*, there is a total sequential order on operations within a single unit of execution, for *co*, there is a total apparent order on all synchronisation operations consistent with *po* and *hb* must be consistent with each *co* and with all *SC* orders.

We illustrate these orderings using the same litmus test.

B. LOADs and STOREs Observed in Program Order (*po*)

Memory accesses from the same master are observed in the order they are issued by all masters.

Figure 3 illustrates *po* on masters *P2* and *P3*.

C. Coherent Observation of LOADs/STOREs per Master (*co*)

If a master observes a STORE followed by a LOAD for the same memory address, then the value it reads in the LOAD is the value defined in the STORE.

Figure 4 illustrates *co* between masters *P2* and *P3*.

D. Happens Before ordering of Load observations (*hb*)

Finally, we introduce the *Happens Before* ordering. This is an ordering between loads defined by a combination of coherent orderings on loads and stores to the same address, i.e., if a load *L1* from an address *A* is observed before a store *S1* to *A* and load *L2* is observed after *S1* by some master, then *L1* is observed before *L2* by all masters. (See 7.)

| P1 | P2 | P3 |
|---------------|----------------|----------------|
| p1i1: S 1, a1 | p2i1: L R1, a2 | p3i1: L R1, a1 |
| p1i2: S 1, a2 | p2i2: L R2, a1 | p3i2: L R2, a2 |

Fig. 7. Happens Before (hb)

E. Synchronisation Mechanisms for Weakly Consistent Memory

Using these notions of observation ordering, *po*, *co*, and *hb*, we can now define the mechanisms by which synchronisation can be achieved. First, we look at the case of *No Synchronisation* for weakly consistent memory

1) *No Synchronisation*: If no fences or specific synchronising instructions are used, then the only ordering is *co*, as shown in Figure 2. No synchronisation means that stores are observed by different masters in different order.

2) *Synchronising Memory Accesses (Single Memory Location)*: HSA supports atomic synchronizing operations with *acquire* and *release* semantics and are, by definition, *Sequentially Consistent*. The operations apply to a single memory location. These synchronising LOADs and STOREs, represented as *SL* and *SS* in Figure 8, result in *po*, *co* and *hb* observation orderings. *SL* and *SS* instructions impose *po* and *hb* ordering.

| P1 | P2 | P3 |
|----------------|-----------------|-----------------|
| p1i1: SS 1, a1 | p2i1: SL R1, a2 | p3i1: SL R1, a1 |
| p1i2: SS 1, a2 | p2i2: SL R2, a1 | p3i2: SL R2, a2 |

Fig. 8. Synchronising Memory Accesses

3) *Fences (Multiple Memory Locations in the same Scope)*: Instead of using synchronising memory accesses, we can introduce a *FENCE* instruction between the LOADs on *P3* and *P4* as shown in Figure 9. The introduction of the *FENCE* also results in *po*, *co* and *hb* observation orderings. Whereas the *LOAD* and *STORE* atomics only apply to a single memory location, the *FENCE* applies to all memory locations within the same defined scope, as described in the HSA specification.

| P1 | P2 | P3 |
|---------------|----------------|----------------|
| p1i1: S 1, a1 | p2i1: L R1, a2 | p3i1: L R1, a1 |
| p1i2: S 1, a2 | p2i2: FENCE | p3i2: FENCE |
| | p2i3: L R2, a1 | p3i3: L R2, a2 |

Fig. 9. Fence Synchronisation

IV. MODELLING WEAK MEMORY WITH EVENT-B

We begin with an abstract, highly nondeterministic model of memory accesses. The abstract Event-B context defines a set of transactions, *TRN* and a subset of the transactions which are memory accesses, *MEMACCESS*.

$axm1 : MEMACCESS \subseteq TRN$

We assume a set of masters that issue transactions on a shared memory space and the effect of those transactions is

observed some time after they are issued. Key to defining weak memory models is the relative ordering, locally and globally, of transaction issues and observations. For this reason, we start our modelling with the key events in weak memory systems, *Issue* transaction and *Observe* transaction.

We therefore define an abstract Event-B machine, *GMI*, with two events, *Issue* and *Observe*, which operate on memory accesses.

```

inv1 : issued  $\subseteq$  MEMACCESS
inv2 : observed  $\subseteq$  issued
Event Issue  $\hat{=}$ 
  any
  where
    ma
  then
    grd1 : ma  $\in$  MEMACCESS
    grd2 : ma  $\notin$  issued
  then
    act1 : issued := issued  $\cup$  {ma}
  end
Event Observe  $\hat{=}$ 
  any
  where
    ma
  then
    grd1 : ma  $\in$  issued
  then
    act1 : observed := observed  $\cup$  {ma}
  end

```

The variable *issued*, invariant *inv1*, initialised to the empty set, records the issuing of each memory access by the event *Issue*. The parameter *ma* of the event *Issue* represents a memory access, the first guard, *grd1*, ensures that the memory access cannot be re-issued if it has already been issued and *ma* is added to the *issued* set by the action *act1*.

The variable *observed*, invariant *inv2* is defined as a subset of the memory accesses that have been issued. If the set *issued* is not empty, the event *Observe* chooses an issued memory access non-deterministically from the set and the action *act1* adds it to the set *observed*. At this early stage there is no notion of a program. Memory accesses are issued in some arbitrary order, and once issued they can be observed *repeatedly* in some arbitrary order.

We then refine the abstract model to differentiate between *LOAD* and *STORE* accesses,

$axm2 : partition(MEMACCESS, STORE, LOAD)$
and refine the abstract events *Issue* and *Observe*.

- New events *IssueLoad* and *IssueStore* refine *Issue*
- New events *ObserveLoad* and *ObserveStore* refine *Observe*

In the next refinement, we introduce an abstraction of the memory architecture, which makes reasoning about weak memory consistency tractable[8]. Each master keeps track individually of its last observed *STORE* value for each memory

location. When a LOAD is observed, the last observed value for that master is written to the master's target register.

An observation ordering is then established between LOADs and STOREs, depending on address, control and data dependencies and the positioning of fences.

The abstract architecture is shown in Figure 10, where each *Memory i* represents the last observed STORE value of each memory location for each individual *Master i*.




Fig. 10. Abstract Memory Architecture

This architecture is represented in Event-B with the variable *lov*.

$inv1 : lov \in MASTER \rightarrow (ADDR \rightarrow VALUE)$
which maintains the *last observed value*.

We also introduce the variable *observers*, which associates each memory access with the masters which have observed that access.

$inv2 : observers \in MEMACCESS \leftrightarrow \mathbb{P}(MASTER)$

For each LOAD or STORE that is issued, *observers(s)* is initialised to the empty set.

The refined event *ObserveStore* introduces a new guard, *grd3*, to ensure that this master, *m*, has not already observed this store, *s*, an action, *act2*, which updates the last observed value for *m*, using the *relational override* operator, and an action, *act3*, which adds *m* to the observers of *s*.

Event *ObserveStore* $\hat{=}$
refines *ObserveStore*

any

s, m

where

grd1 : *s* \in *issued*

grd2 : *s* \in *STORE*

grd3 : *m* \notin *observers(s)*

then

act1 : *observed* := *observed* \cup {*s*}

act2 : *lov(m)* := *lov(m)* \Leftarrow {*addr(s)* \mapsto *val(s)*}

act3 : *observers(s)* := *observers(s)* \cup {*m*}

end

We now have a generic, abstract model which can now be refined to implement HSA Atomics and Fences.

A. Modelling HSA Fence Synchronisation

Using the definitions of *po*, *co* and *hb* from the HSA specification, we are able to model the observation orderings

necessary to implement fence synchronisation. We introduce the variable *ahead* to represent the accesses that are ahead of the fence and the variable *after* to represent the set of LOADs that are observed after a given STORE. Our refinement distinguishes multiple cases of LOAD observation, each of which is represented by an event that refines our abstract *ObserveLoad* event. First, we specify the LOAD observations in the presence of a FENCE. There are two cases to address, the load happening before a store and happening after a store.

ObserveLoadHappensBeforeWithFence: If the master that issued the LOAD has issued a FENCE and if the LOAD is not ahead of the FENCE in program order, then all other memory accesses ahead of the FENCE in program order must have been observed (guard *grd7* below). This LOAD is observed before a corresponding STORE to the same location (guards *grd9*, *grd10*) and no other LOAD has observed that STORE and there is therefore no LOAD *after* the STORE (guard *grd11*).

- LOADs are observed in *Program Order(po)*
- LOAD is observed before a corresponding STORE
- No other LOAD has been observed after the STORE
- *po* + *co* + *hb*

The event is specified as follows:

Event *ObserveLoadHappensBeforeWithFence* $\hat{=}$
extends *ObserveLoad*

any

l, m, f, s

where

grd1 : *l* \in *issued*

grd2 : *l* \in *LOAD*

grd3 : *m* \notin *observers(l)*

grd4 : *m* = *issuer(l)*

grd5 : *f* \in *issuedfence*

grd6 : *m* = *issuer(f)*

grd7 : *l* \notin *ahead(f)* \Rightarrow

($\forall a. a \in \text{ahead}(f) \wedge a \in \text{dom}(\text{observers}) \Rightarrow$
m \in *observers(a)*)

grd8 : *s* \in *STORE*

grd9 : *address(s)* = *address(l)*

grd10 : *m* \notin *observers(s)*

grd11 : *after(s)* = \emptyset

then

act1 : *observed* := *observed* \cup {*l*}

act2 : *observers(l)* := *observers(l)* \cup {*m*}

end

ObserveLoadAfterStoreWithFence: If the master that issued the LOAD has issued a FENCE and if the LOAD is not ahead of the FENCE in program order, then all other memory accesses ahead of the FENCE in program order must have been observed (guard *grd7*). The master has observed a corresponding STORE to the same location (guard *grd11*).

- In Program Order
- LOAD is observed after a corresponding STORE
- *po* + *co*

The event is specified as follows:

Event *ObserveLoadAfterStoreWithFence* $\hat{=}$
extends *ObserveLoad*
any
 l, m, f, s
where
 grd1 : $l \in \text{issued}$
 grd2 : $l \in \text{LOAD}$
 grd3 : $m \notin \text{observers}(l)$
 grd4 : $m = \text{issuer}(l)$
 grd5 : $f \in \text{issuedfence}$
 grd6 : $m = \text{issuer}(f)$
 grd7 : $l \notin \text{ahead}(f) \Rightarrow$
 $(\forall a. a \in \text{ahead}(f) \wedge a \in \text{dom}(\text{observers}) \Rightarrow$
 $m \in \text{observers}(a))$
 grd8 : $s \in \text{issued}$
 grd9 : $s \in \text{STORE}$
 grd10 : $\text{address}(s) = \text{address}(l)$
 grd11 : $m \in \text{observers}(s)$
then
 act1 : $\text{observed} := \text{observed} \cup \{l\}$
 act2 : $\text{observers}(l) := \text{observers}(l) \cup \{m\}$
 act3 : $\text{after}(s) := \text{after}(s) \cup \{l\}$
end

Similarly, we specify the two cases of LOAD observation *without* a FENCE covering loads before and after a store. We also specify an *IssueFence* event. Details of these are straightforward and are omitted for brevity.

B. Modelling the HSA Atomics

Again, using the definitions of *po*, *co* and *hb* from the HSA specification, we are now able to model the HSA atomics.

First, we define the atomic LOAD, *SCACQLOAD*, with *acquire* semantics and the atomic STORE, *SCRELSTORE* with *release* semantics, and differentiate them from the *ordinary* LOADs and STOREs.

As we did to model fences, above, we then refine the events of the generic model to establish this differentiation, imposing the ordering defined in the HSA specification in the following, refined *Issue* events, *IssueAtomicSCACQLoad*, *IssueLoad*, *IssueAtomicSCRELStore*, *IssueStore* and their corresponding *Observe* events.

C. Introducing the Register File

We have a model which implements fences and atomics, which can be refined to introduce the register file as an Event-B variable, *rf*

inv1 : $rf \in \text{MASTER} \rightarrow (\text{REG} \rightarrow \text{VALUE})$

When a LOAD, *l*, is observed by a master, *m*, the register *r* associated with *l*, *rl*, takes the last observed value for that LOAD location, *addr(l)*. For instance, in the event *ObserveLoadHappensBeforewithFence* described above, an extra action, *act3* is added.

act3 : $rf(m) := rf(m) \Leftarrow \{r(l) \mapsto (\text{lov}(m))(\text{addr}(l))\}$

V. VALIDATING THE MEMORY MODEL

A. A Notation for Representing Litmus Tests

The memory model we have developed represents all the legal interleavings of an arbitrary set of programs running on an arbitrary set of masters. We now wish to constrain the model to represent a litmus test - all possible interleavings for a fixed set of masters each running a specific, fixed program thread.

To make it easier to verify litmus tests, we introduce a tool-supported notation with which the litmus test and the invariant can be represented and an Event-B *context* and *invariant* is generated automatically from this description. An example of a litmus test in our notation as shown in Figure 11.

```

p1:
  st 1, a1
  st 1, a2

p2:
  ld r1, a1
  fence
  ld r2, a2

p3:
  ld r1, a2
  fence
  ld r2, a1

outcome: not(p2.r1 = 1 and p3.r1 = 1 and p2.r2 = 0 and p3.r2 = 0)

```

Fig. 11. Litmus Test in Tool-supported Notation

B. The Litmus Test Context

We constrain the memory model by introducing axioms which specify the litmus test instructions, where for instance *I11* is a STORE of the value *1* to the address *a1*, and the program threads that will run those instructions.

```

axm10 :  $P1 = \{1 \mapsto I11, 2 \mapsto I12\}$ 
axm11 :  $P2 = \{1 \mapsto I21, 2 \mapsto I22, 3 \mapsto I23\}$ 
axm12 :  $P3 = \{1 \mapsto I31, 2 \mapsto I32, 3 \mapsto I33\}$ 

```

We then associate each thread with the appropriate master.

```

axm19 :  $\text{PROGRAM} = \{M1 \mapsto P1, M2 \mapsto P2, M3 \mapsto P3\}$ 

```

C. Verifying the Litmus Test Outcome

An *outcome* statement may be associated with the Litmus Test, which is represented as a boolean expression in terms of the values in the registers of the masters. We specify *outcome* to mean *the values in the registers when all litmus test loads have been observed*. An invariant is generated that we verify using the ProB model checker. In the case of this litmus test, the invariant is

```

inv1 :  $\{I21, I23, I31, I33\} \subseteq \text{observed} \Rightarrow$ 
 $\neg((rf(M2))(R1) = V1 \wedge (rf(M3))(R1) = V1 \wedge$ 
 $((rf(M2))(R2) = V0 \wedge (rf(M3))(R2) = V0))$ 

```

If the loads, *I21*, *I23*, *I31*, *I33* issued by the masters *M2* and *M3* have been observed, then the registers associated with the masters cannot contain the prohibited outcome. The values *V0* and *V1* represent logic 0 and 1 respectively.

D. Covering the Allowable Register Value Combinations

The facility to specify and verify the illegal litmus test outcomes formally is valuable, but to ensure thorough verification, it is necessary to go further.

- To ensure that all the *legal outcomes* of the litmus test are *reachable*.
- To ensure that the set of litmus tests *cover* the model functionality.

We introduce a further refinement of the memory model to introduce a *functional coverage metric* for the litmus test.

The relation *coverage*

$inv1 : coverage \in REGCOVER \leftrightarrow REGCOVER$ represents the combinations of register values that are reached by the litmus test, where *REGCOVER* is defined in extended context thus:

```
axm1 : C0 = {R1 ↦ V0, R2 ↦ V0}
axm2 : C1 = {R1 ↦ V0, R2 ↦ V1}
axm3 : C2 = {R1 ↦ V1, R2 ↦ V0}
axm4 : C3 = {R1 ↦ V1, R2 ↦ V1}
axm5 : REGCOVER = {C0, C1, C2, C3}
```

Each constant C_n represents a register combination value for a master and *REGCOVER* is the set of all register combinations for a master. The *coverage* relation is initialised to the empty set. The refinement then introduces a new event *CoverRegisterValues*

Event *CoverRegisterValues* $\hat{=}$

any

$rm2, rm3$

where

```
grd1 : rm2 ∈ REGCOVER
grd2 : rm3 ∈ REGCOVER
grd3 : rf(M2) = rm2
grd4 : rf(M3) = rm3
grd5 : rm2 ↦ rm3 ∉ coverage
grd6 : {I21, I23, I31, I33} ⊆ observed
```

then

```
act1 : coverage := coverage ∪ {rm3 ↦ rm4}
```

end

This event is enabled when all the LOADs have been observed (guard *grd6*) and when this combination of the master register values has not already been covered (guard *grd5*). The combination is then added to the coverage relation.

We can now re-run the model checker exhaustively to show, not only that illegal register combinations are never reached but also that all legal combinations are reachable. The ProB model checker has a coverage option that now allows us to display the values of the relation *coverage* that are covered.

Coverage point 15 represents the initial value of the coverage relation. Note that otherwise all 15 possible combinations have been covered and the illegal combination is not covered.

E. Generating a Test for each Coverage Point

Now that we are sure that each legal combination of registers is reachable, we wish to create a regression test

for each of the 15 reachable combinations. Creating the tests manually is time consuming. ProB provides a *model-checking-based* test generation facility which can be used to generate these tests automatically.

First, we specify a predicate to define the target state. For instance, if we wish to show that all the registers of both masters can take the value 0 when all the LOADs have been observed, then we specify.

$$I21, I23, I31, I33 \subseteq observed \Rightarrow C0 \mapsto C0 : coverage$$

Second, we specify the event or events which we wish to be covered by the test. In this case, we just specify the *CoverRegisterValues* event to guide the search.

ProB generates the tests as an HTML file, which can be translated to suit the tool chain. One of the tests that ProB generates is shown in Figure 12 below. The test describes a sequence of six *Issue* instructions which results in the expected register value combination to be recorded by the event *CoverRegisterValues*.

```
Test_case_id = 1
IssueLoad(I21, P2);
IssueLoad(I31, P3);
IssueFence(I22, P2);
IssueLoad(I23, P2);
IssueFence(I32, P3);
IssueLoad(I33, P3);
CoverRegisterValues(rm2(R1->0, R2->0),
                    rm3(R1->0, R2->0));
```

Fig. 12. Generated Test

F. Measuring the Coverage of a set of Litmus Tests

Since the litmus tests each target a particular aspect of the weak memory model, an individual litmus test will not cover all of the events of the model. What we do want to know, however, is that all events of the model are covered by the set of litmus tests. In other words, that the litmus tests *fully cover* the model functionality.

We return to the tests generated above and now, instead of specifying the *CoverRegisterValues* event, we specify *all* the events that we expect to be covered by the test. We want to verify that these and only these events are covered by the litmus test. ProB will now generate several tests for the required outcome which cover the events specified.

G. Generating additional Litmus Tests

Once we have completed the model validation and have a set of litmus tests which cover the model, we relax the constraints on our formalised litmus tests so that each model represents a class of programs with a non-deterministic mix of instructions (and the same set of masters, registers and address space as the associated litmus). Using the ProB test generator, we are able to generate automatically a wide range of tests, with expected

register values, which can be used for regression testing of the HSA-compliant platform itself.

VI. RELATED WORK AND CONCLUSION

An overview of memory ordering and barriers in modern microprocessors is presented in [9] and [10]. The release consistency model is described in [11]. The notion of using thread-local re-ordering and store atomicity, together with the definition of an abstract architecture to support weakly consistent memory specification is developed in [8] and [12]. The tutorial [13], for ARM and POWER memory models, builds on the notion of an abstract machine architecture to define a *storage subsystem* which has general applicability for weakly consistent memory specification. [14] introduces the notion of modelling concurrency with partial orders. [15] presents an axiomatic approach to POWER memory modelling and [16] presents a class of relaxed memory models which are parameterised to support different local re-orderings and store atomicity relaxation. [6] provides a clear description of the ARM memory model and the HSA initiative to harmonise heterogeneous computing concepts [2] has resulted in a natural language specification of fences and atomics which provides considerable clarity for programmers, model and system developers.

In [17], an *axiomatic* approach is used where the memory model is represented as a directed graph, the nodes representing the memory accesses and the arcs the ordering between those accesses. Our Event-B approach is based on the computation model of *guarded atomic actions*. Memory access ordering is specified *axiomatically* as guards on events, which model the memory accesses themselves, for instance *IssueRead*, *ObserveWrite*. If more than one event is enabled, then a choice is made non-deterministically, naturally representing the non-determinism of the memory model. The Event-B model is operational in the sense that it is based on the highly abstract memory architecture described in [8]. We begin with a generic model to represent the ordering at a high level of abstraction and then use formal refinement to introduce, in steps, as much detail as is necessary to represent a concrete, concurrent program, where the state of the registers and the values observed by each of the masters are updated by the event actions. The concrete model can therefore be mapped directly to a Transaction-Level Model (TLM) for efficient system simulation. [17] also presents a tool for simulation and testing of memory models. Using ProB we can simulate and test our Event-B models at each level of abstraction as well as generate coverage-driven tests for regression testing of hardware and software systems.

We believe our methodology of refinement based formal modelling, baseline compliance testing of the formal model and coverage driven test extraction using a single language of Event B language and the Rodin tool is a completely new way of addressing a profoundly important challenge facing the design and verification of low-power of multi-core systems. It provides a hierarchy of models where the complexity of modelling is managed by refinement leading to clarity

in understanding and precision of formal modelling. Using the same Event B model on which we have shown that refinements are provably correct, we can test it for litmus test compliance against published HSA litmus tests. The key here is that testing for compliance is embedded within the Rodin framework where proof based refinement is done; leading to a very tight integration of testing and proofs. By relaxing the constraints of HSA litmus tests we are able to explore a much bigger space of tests and by a natural extension of the core Event B model of the HSA by a coverage event we are able to generate a bigger set of regression tests that satisfy the coverage model. This is again done on the same core Event B model on which proof based refinements were done. The tests generated can be used in the development and verification of the HSA-compliant platform itself.

ACKNOWLEDGMENT

The authors would like to thank Jason Meredith, Mark Landers and James Aldis for several discussions during the development of this work. The authors would also like to thank Imagination Technologies for their time and support throughout this work.

REFERENCES

- [1] M. Dubois, C. Scheurich, and F. Briggs, "Memory access buffering in multiprocessors," *ACM SIGARCH Computer Architecture News*, vol. 14, no. 2, pp. 434–442, 1986.
- [2] H. Foundation, "HSA platform system architecture specification," Standard, January 2015.
- [3] J.-R. Abrial, M. J. Butler, S. Hallerstede, T. S. Hoang, F. Mehta, and L. Voisin, "Rodin: an open toolset for modelling and reasoning in Event-B," *STTT*, vol. 12, no. 6, pp. 447–466, 2010.
- [4] L. Lamport, "How to make a multiprocessor computer that correctly executes multiprocess programs," *Computers, IEEE Transactions on*, vol. 100, no. 9, pp. 690–691, 1979.
- [5] P. Sewell, S. Sarkar, S. Owens, F. Z. Nardelli, and M. O. Myreen, "x86-tso: a rigorous and usable programmer's model for x86 multiprocessors," *Communications of the ACM*, vol. 53, no. 7, pp. 89–97, 2010.
- [6] N. Chong and S. Ishtiaq, "Reasoning about the arm weakly consistent memory model," in *Proceedings of the 2008 ACM SIGPLAN workshop on Memory systems performance and correctness: held in conjunction with the Thirteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'08)*. ACM, 2008, pp. 16–19.
- [7] M. Leuschel and M. Butler, "ProB: an automated analysis toolset for the B method," *STTT*, vol. 10, no. 2, pp. 185–203, 2008.
- [8] X. Shen, Arvind, X. Shen, L. Rudolph, and L. Rudolph, "Commit-reconcile fences (crf): A new memory model for architects and compiler writers," in *In Proceedings of the 26th International Symposium on Computer Architecture*, 1999, pp. 150–161.
- [9] P. E. McKenney, "Memory ordering in modern microprocessors, part i," *Linux Journal*, vol. 2005, no. 136, p. 2, 2005.
- [10] —, "Memory ordering in modern microprocessors, part ii," *Linux Journal*, vol. 2005, no. 137, p. 5, 2005.
- [11] K. Gharachorloo, D. Lenoski, J. Laudon, P. Gibbons, A. Gupta, and J. Hennessy, *Memory consistency and event ordering in scalable shared-memory multiprocessors*. ACM, 1990, vol. 18, no. 2SI.
- [12] A. Arvind and J.-W. Maessen, "Memory model= instruction reordering+ store atomicity," in *ACM SIGARCH Computer Architecture News*, vol. 34. IEEE Computer Society, 2006, pp. 29–40.
- [13] L. Maranget, S. Sarkar, and P. Sewell, "A tutorial introduction to the arm and power relaxed memory models," 2012.
- [14] V. Pratt, "Modeling concurrency with partial orders," *International Journal of Parallel Programming*, vol. 15, no. 1, pp. 33–71, 1986.

- [15] S. Mador-Haim, L. Maranget, S. Sarkar, K. Memarian, J. Alglave, S. Owens, R. Alur, M. M. Martin, P. Sewell, and D. Williams, "An axiomatic memory model for power multiprocessors," in *Computer Aided Verification*. Springer, 2012, pp. 495–512.
- [16] J. Alglave, L. Maranget, S. Sarkar, and P. Sewell, "Fences in weak memory models," in *Computer Aided Verification*. Springer, 2010, pp. 258–272.
- [17] J. Alglave, L. Maranget, and M. Tautschnig, "Herding cats: Modelling, simulation, testing, and data mining for weak memory," *ACM Transactions on Programming Languages and Systems (TOPLAS)*, vol. 36, no. 2, p. 7, 2014.